

AD-A181 026 SHOULD THE AIR FORCE PERSONNEL DATA SYSTEM USE DATABASE 1/1
MACHINES?(U) AIR COMMAND AND STAFF COLL MAXWELL AFB AL
M T ANDERSON APR 87 ACSC-87-0110

AD-A181 026 SHOULD THE AIR FORCE PERSONNEL DATA SYSTEM USE DATABASE 1/1
MACHINES?(U) AIR COMMAND AND STAFF COLL MAXWELL AFB AL
M T ANDERSON APR 87 ACSC-87-0110


AD-A181 026 SHOULD THE AIR FORCE PERSONNEL DATA SYSTEM USE DATABASE 1/1
MACHINES?(U) AIR COMMAND AND STAFF COLL MAXWELL AFB AL
M T ANDERSON APR 87 ACSC-87-0110

UNCLASSIFIED H I ANDERSON APR 87 ACSC-87-0110 F/G 12/7 NL

UNCLASSIFIED H I ANDERSON APR 87 ACSC-87-0110 F/G 12/7 NL

UNCLASSIFIED H I ANDERSON APR 87 ACSC-87-0110 F/G 12/7 NL

[illegible]



AD-A181 026



AIR COMMAND AND STAFF COLLEGE

STUDENT REPORT

SHOULD THE AIR FORCE PERSONNEL
DATA SYSTEM USE DATABASE MACHINES?

MAJOR MICHAEL T. ANDERSON

87-0110

"insights into tomorrow"

DTIC
ELECTE

JUN 11 1987

04
E

This document has been approved
for release and distribution
by the DTIC/STC

DISCLAIMER

The views and conclusions expressed in this document are those of the author. They are not intended and should not be thought to represent official ideas, attitudes, or policies of any agency of the United States Government. The author has not had special access to official information or ideas and has employed only open-source material available to any writer on this subject.

This document is the property of the United States Government. It is available for distribution to the general public. A loan copy of the document may be obtained from the Air University Interlibrary Loan Service (AUL/LDEX, Maxwell AFB, Alabama, 36112) or the Defense Technical Information Center. Request must include the author's name and complete title of the study.

This document may be reproduced for use in other research reports or educational pursuits contingent upon the following stipulations:

- Reproduction rights do not extend to any copyrighted material that may be contained in the research report.

- All reproduced copies must contain the following credit line: "Reprinted by permission of the Air Command and Staff College."

- All reproduced copies must contain the name(s) of the report's author(s).

- If format modification is necessary to better serve the user's needs, adjustments may be made to this report--this authorization does not extend to copyrighted information or material. The following statement must accompany the modified document: "Adapted from Air Command and Staff Research Report (number) entitled (title) by (author) ."

- This notice must be included with any reproduced or adapted portions of this document.



REPORT NUMBER 87-0110

TITLE SHOULD THE AIR FORCE PERSONNEL DATA SYSTEM USE DATABASE MACHINES?

AUTHOR(S) MAJOR MICHAEL T. ANDERSON, USAF

FACULTY ADVISOR MAJOR TERRY L. BROOKS, ACSC 3823 STUS

SPONSOR MR DAVID ESSENPREIS, AFMPC/DPMDX

Submitted to the faculty in partial fulfillment of
requirements for graduation.

AIR COMMAND AND STAFF COLLEGE
AIR UNIVERSITY
MAXWELL AFB, AL 36112

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE

REPORT DOCUMENTATION PAGE

1a. REPORT SECURITY CLASSIFICATION UNCLASSIFIED		1b. RESTRICTIVE MARKINGS										
2a. SECURITY CLASSIFICATION AUTHORITY		3. DISTRIBUTION/AVAILABILITY OF REPORT STATEMENT "A" Approved for public release; Distribution is unlimited.										
2b. DECLASSIFICATION/DOWNGRADING SCHEDULE												
4. PERFORMING ORGANIZATION REPORT NUMBER(S) 87-0110		5. MONITORING ORGANIZATION REPORT NUMBER(S)										
6a. NAME OF PERFORMING ORGANIZATION ACSC/EDCC	6b. OFFICE SYMBOL (If applicable)	7a. NAME OF MONITORING ORGANIZATION										
6c. ADDRESS (City, State and ZIP Code) Maxwell AFB AL 36112-5542		7b. ADDRESS (City, State and ZIP Code)										
8a. NAME OF FUNDING/SPONSORING ORGANIZATION	8b. OFFICE SYMBOL (If applicable)	9. PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER										
8c. ADDRESS (City, State and ZIP Code)		10. SOURCE OF FUNDING NOS. <table border="1"><tr><td>PROGRAM ELEMENT NO.</td><td>PROJECT NO.</td><td>TASK NO.</td><td>WORK UNIT NO.</td></tr></table>		PROGRAM ELEMENT NO.	PROJECT NO.	TASK NO.	WORK UNIT NO.					
PROGRAM ELEMENT NO.	PROJECT NO.	TASK NO.	WORK UNIT NO.									
11. TITLE (Include Security Classification) SHOULD THE AIR FORCE PERSONNEL												
12. PERSONAL AUTHOR(S) Anderson, Michael T., Major, USAF												
13a. TYPE OF REPORT	13b. TIME COVERED FROM _____ TO _____	14. DATE OF REPORT (Yr., Mo., Day) 1987 April	15. PAGE COUNT 35									
16. SUPPLEMENTARY NOTATION ITEM 11: DATA SYSTEM USE DATA BASE MACHINES?												
17. COSATI CODES <table border="1"><tr><td>FIELD</td><td>GROUP</td><td>SUB. GR.</td></tr><tr><td></td><td></td><td></td></tr><tr><td></td><td></td><td></td></tr></table>		FIELD	GROUP	SUB. GR.							18. SUBJECT TERMS (Continue on reverse if necessary and identify by block number)	
FIELD	GROUP	SUB. GR.										
19. ABSTRACT (Continue on reverse if necessary and identify by block number) Database management systems are in wide use today for large automated information systems. Special-purpose computers, called database machines, are becoming commercially available. These computers are tailored to manipulate structured databases very efficiently. This study examines the suitability of one such machine for use in the Air Force Personnel Data System (PDS). The study concludes that the machine, the Teradata DBC/1012, has the potential to improve the performance of the PDS, and should be considered for purchase.												
20. DISTRIBUTION/AVAILABILITY OF ABSTRACT UNCLASSIFIED/UNLIMITED <input type="checkbox"/> SAME AS RPT. <input checked="" type="checkbox"/> DTIC USERS <input type="checkbox"/>		21. ABSTRACT SECURITY CLASSIFICATION UNCLASSIFIED										
22a. NAME OF RESPONSIBLE INDIVIDUAL ACSC/EDCC Maxwell AFB AL 36112-5542		22b. TELEPHONE NUMBER (Include Area Code) (205) 293-2483	22c. OFFICE SYMBOL									

PREFACE

The past twenty years have seen significant improvements in the computer technology available for managing large amounts of information, progressing from simple file systems in the 1960s to the integrated database management systems of today. For some time, researchers have been studying special-purpose computers designed for the efficient manipulation of large databases. Such machines are now becoming commercially available. This study is an initial look at the advantages of using one of these "database machines" to improve the performance of the Air Force Personnel Data System.

Accession For	
NRIS GRA&I	<input checked="checked" type="checkbox"/>
LTIC TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	
By	
Distribution/	
Availability Codes	
Avail and/or	
Dist	Special
A-1	



ABOUT THE AUTHOR

Major Michael T. Anderson has had a varied Air Force career. He enlisted in 1972, immediately after graduating from college, and served for two years as an administration specialist in an early-warning radar detachment. He was commissioned through Officer Training School in 1974. He was assigned to the NORAD Cheyenne Mountain Complex as a space systems officer, where he worked in the missile warning center until 1976, when he was transferred to Clear AFS, Alaska, to operate the Ballistic Missile Early Warning System radar. In 1977, he moved to Headquarters, Strategic Air Command, and cross-trained into the computer systems career field. While at Hq SAC, Major Anderson developed several new intelligence production methods for the SAC Intelligence Data Handling System. After some time out for school assignments (AFIT Master's program and Squadron Officer School), he was assigned to the Air Force Military Personnel Center (AFMPC) as a personnel data systems officer until August 1986. For the last three years of his tour at AFMPC, Major Anderson served as the chief database administrator and designer for the Air Force Personnel Data System (PDS), and was instrumental in the successful migration of the headquarters-level PDS between two different makes of computers.

Major Anderson holds a Bachelor of Science in Physics from the University of California, Riverside; a Master of Science in Computer Engineering from Stanford University; and is a distinguished graduate of the Air Force Squadron Officer School. In addition to his military experience and academic education, he has taught computer science for several years at both the graduate and undergraduate levels, at St. Mary's University, San Antonio, Texas.

TABLE OF CONTENTS

Preface.....	iii
List of Illustrations.....	vi
Executive Summary.....	vii
CHAPTER 1 - THE PROBLEM	
Workload of the Air Force Personnel Data System.....	1
The Shortfalls of REACQ.....	1
A Possible Solution.....	2
CHAPTER 2 - THE AIR FORCE PERSONNEL DATA SYSTEM	
Missions.....	3
System Description.....	3
CHAPTER 3 - TECHNOLOGICAL BACKGROUND	
Introduction.....	5
Database Management Systems.....	5
Database Machines.....	11
CHAPTER 4 - COMPARISON OF DATABASE MACHINE AGAINST A CONVENTIONAL DATABASE MANAGEMENT SYSTEM	
Selecting a Database Machine for the PDS.....	15
Qualitative Factors.....	15
Quantitative Factors.....	17
CHAPTER 5 - CONCLUSIONS, FINDINGS, AND RECOMMENDATIONS	
Introduction.....	21
Conclusions.....	21
Findings.....	22
Recommendations.....	23
BIBLIOGRAPHY.....	25

LIST OF ILLUSTRATIONS

TABLES

TABLE 1--Commercial Database Machines.....	13
TABLE 2--AA Performance Estimates.....	20
TABLE 3--Prices of Teradata DBC/1012 Systems.....	20

FIGURES

FIGURE 1--Hierarchical Database Model.....	6
FIGURE 2--Network Database Model.....	8
FIGURE 3--Relational Database Model.....	10
FIGURE 4--Teradata DBC/1012 Configuration.....	16
FIGURE 5--Teradata DBC/1012 Performance Curves.....	19



EXECUTIVE SUMMARY

Part of our College mission is distribution of the students' problem solving products to DoD sponsors and other interested agencies to enhance insight into contemporary, defense related issues. While the College has accepted this product as meeting academic requirements for graduation, the views and opinions expressed or implied are solely those of the author and should not be construed as carrying official sanction.

"insights into tomorrow"

REPORT NUMBER 87-0110

AUTHOR(S) MAJOR MICHAEL T. ANDERSON. USAF

TITLE SHOULD THE AIR FORCE PERSONNEL DATA SYSTEM USE
DATABASE MACHINES?

I. Purpose: To determine whether the performance, cost, and reliability of the Air Force Personnel Data System can be significantly improved by installing a commercially available special-purpose computer dedicated to database management tasks (a "database machine").

II. Problem: The PDS has, for years, been constrained in the services it can provide, primarily due to performance limitations imposed by its computer systems. The hardware reacquisition project (REACQ) of 1983-85 was an attempt to deal with this problem, but was only partially successful. A new approach to data processing may offer a breakthrough in the performance level of the PDS, thereby making personnel data services cheaper, more reliable, and more available.

III. Data: The PDS includes a wide range of applications, ranging from traditional record-keeping to recruiting to force structure modeling to formal training management, and much more. Utilization statistics show, however, that nearly 80% of the entire central-site computer resource is dedicated to only a few general categories: the master personnel files, and a few, very active, online information systems (PMS, PROMIS, and SURF). These systems are very much "input-output intensive," spending most of their time reading or writing data to disk devices. Furthermore, all of these systems use a database management system

CONTINUED

(DBMS) to handle their I/O operations. Any effort to improve the performance of the PDS must address these systems, and must certainly address the performance of the underlying DBMS. Traditional database architectures have been tailored for existing mainframe computers in what has been described as a "machine-friendly" design. These architectures are geared for "record-at-a-time" processing, and use intricate chain or tree structures to speed up the process of finding data. Recently, the more "user-friendly" relational database architecture has become commercially available in packages which are reasonably efficient for small-to-medium sized applications, but this "elegant" approach to database management has been too computationally intensive to be successful for large information systems with many users. Special-purpose computers, called "database computers" or "database machines," have recently appeared on the commercial market. These computers offer the capability of applying sufficient computational power to the needs of a relational database architecture to make its performance comparable to, or even better than, traditional database management systems. As of this writing, only one commercially available database machine has the storage capacity and hardware compatibility necessary for the PDS; that machine is the Teradata DBC/1012. A case-study comparison of potential performance, rating the DBC/1012 against the current Honeywell DPS-8/70 DM-IV system, shows that the DBC/1012 offers several distinct advantages. The DBC/1012 has the potential for improved performance, simpler backup and recovery operations, less susceptibility to failure, and a greater capacity for growth.

IV. Conclusions: The Teradata DBC/1012 database machine compares favorably with the mainframe DBMS operating at AFMPC. It offers improvements in system performance, ease of operation, reliability, and ease of software development.

V. Recommendation: The Director of Personnel Data Systems, Air Force Military Personnel Center, should begin action to develop prototype relational database software for the purpose of benchmarking the Teradata DBC/1012 database computer. If the benchmark demonstrates a significant performance improvement over the existing mainframe DBMS, the Director should take action to procure a DBC/1012 large enough to support the database portion of the PDS.

Chapter 1

THE PROBLEM

WORKLOAD OF THE AIR FORCE PERSONNEL DATA SYSTEM

The US Air Force Personnel Data System (PDS), operated from the central site at the Air Force Military Personnel Center (AFMPC), Randolph AFB, Texas, is the largest personnel data system in the federal government. It operates 24 hours a day, seven days a week, managing the personnel records of over one million active duty, Air Reserve Forces, and civilian personnel. It supports over 150 Air Force bases worldwide, as well as some fifty major air commands, separate operating agencies, and intermediate headquarters. Since the mid-1970s, the PDS has suffered capacity problems, forcing the limitation or curtailment of certain data processing services. In the late 1970s, an initiative to expand the capacity of the PDS was begun, which would improve the system through the purchase of new computer hardware. In the fall of 1982, the Honeywell Corporation was awarded the contract for the AFMPC Reacquisition Project (REACQ). The goal of REACQ was to replace the existing Burroughs 6700 computer complex with newer, high-performance mainframe computers to improve the overall performance and capacity of the PDS.

THE SHORTFALLS OF REACQ

The new computer hardware purchased under REACQ required that most of the software for the PDS be modified to operate on the new computers; this modification was completed in the fall of 1985. In the three years since the initial contract award, the nature and extent of much of the central site software for the PDS had changed, but the fundamental mission and modes of processing remained basically the same. Some data processing applications had grown significantly, or had been greatly improved, others became obsolete or radically altered. Overall, the software supporting the PDS was adapted to fit the new hardware and software environment without any serious deficiencies or loss of capability. However, one major problem remains. As of today, the PDS is again facing a capacity problem, similar to that of the mid-1970s. While the replacement of computers did provide the PDS with new capacity, the modification of old software, and the development of new, has absorbed much of that added capacity. Today, the average user of the PDS in the major air command or Air Staff office sees little

improvement in the overall performance of the PDS as compared to ten years ago.

A POSSIBLE SOLUTION

Given this somewhat dreary history, is the PDS doomed to continuing cycles of "playing catch-up," attempting to keep computer capacity ahead of workload? Traditional computer system capacity planning methods might lead to that scenario. Unfortunately, the Air Force has neither the time nor money to continue operating in this fashion. Fortunately, the advent of several new hardware and software technologies give us the capability to break out of the "catch-up" cycle. One of the promising technologies of the 1980s is the backend database machine. This is a special-purpose computer designed to extend the capacity of a general-purpose computer system and prolong its useful lifetime by adding processing capacity without the need for a complete overhaul of software and a complete replacement of computer hardware at a data processing installation. In this paper I will attempt to explain the advantages of a database machine, and evaluate its usefulness in the operation of the PDS. Based on this evaluation, I will answer the question "Should a database machine be used as an integral part of the Air Force Personnel Data System?"

Chapter 2

THE AIR FORCE PERSONNEL DATA SYSTEM

MISSIONS

The Air Force Military Personnel Center is tasked with operating the overall personnel system for all members of the US Air Force. To do this, the center is authorized to

. . . develop and implement policy concerning accessions testing, classification, worldwide distribution and management of personnel, automated personnel systems, military personnel records systems, standard personnel operations, programs, officer and airman performance evaluation, promotion testing, reenlistment and retention, leave, survivors benefits, escort and dependent travel, awards and decorations, appearance standards, nonappropriated fund manpower requirements, Morale, Welfare, and Recreation (MWR) activities, active duty service commitments, specified period of time contracts, and overseas tour lengths. Assist in the development and implementation of policies pertaining to procurement . . . promotion of officer and enlisted members, demotion of enlisted members, desertion, absent without leave (AWOL), Regular-Reserve-temporary Air Force appointments, separations, retirements, flying status, service dates, Indefinite Reserve Status, and Social Actions programs. (16:1)

A key tool in operating a personnel system dealing with over 1 million people (6:190) is a extensive computer system. AFR 23-33 gives AFMPC the authority for operating, scheduling, and maintaining the central site facility for the Personnel Data Systems computers and peripheral equipment to support the Air Staff, AFMPC, USAFR, ANG, MAJCOMs, separate operating agencies, and base-level consolidated base personnel offices. (16:2)

SYSTEM DESCRIPTION

The Air Force Personnel Data System is a collection of computer systems located at each Air Force base, Headquarters US Air Force, and at the Air Force Military Personnel Center at Randolph Air Force Base, Texas. These systems are interconnected with a variety of

communications systems, including high-speed data links, low-speed telephone lines, and AUTODIN message connections. The hub of the entire personnel data processing effort is the central site at the Military Personnel Center (MPC). The central site at MPC is responsible for the maintenance of a single centralized repository of data used by the entire Air Force. A "master record" is kept for every individual in the active and reserve forces in the master personnel files. The schedules, class rosters, and management information for all Air Force formal training are maintained in the Pipeline Management System (PMS). New Air Force recruits are matched against job requirements and training schedules using the Procurement Management Information System, or PROMIS. The specialized information for airman and officer promotions is kept up-to-date and is used by promotion boards and the Weighted Airman Promotion System. Finally, several other smaller collections of data are maintained for such diverse areas as force structure, officer assignments, Congressional Inquiries, and the Air Force Suggestion Program. (14:1-2 - 1-7)

Chapter 3

TECHONOLOGICAL BACKGROUND

INTRODUCTION

Database management systems (DBMSs) are complex combinations of computer hardware and software. Originally, information on a computer system was simply managed using the vendor-supplied file system, part of the operating system furnished with the computer. Over the years, however, specialized software and hardware have been developed to handle large databases, for a variety of reasons. These range from efforts to increase programmer productivity to providing built-in protection for data checking and recovery. Initially, the emphasis was placed on specialized software, running on general-purpose computers. More recently, special-purpose computers have been designed to use "streamlined" versions of this database management software.

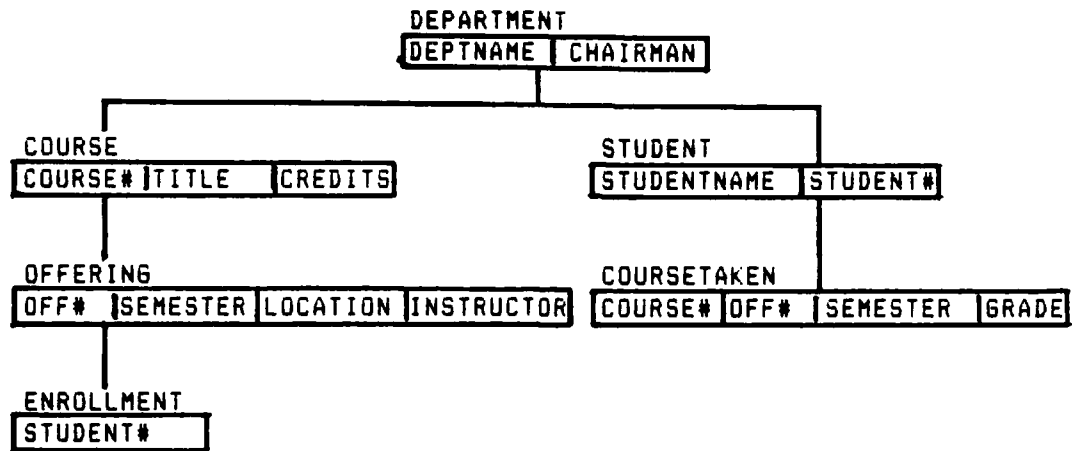
DATABASE MANAGEMENT SYSTEMS

Database management systems have "come into their own" in the last decade, primarily because of reliability and ease of use. Prior to the widespread use of DBMSs in commercial and government applications, programmers used their own "homegrown" file structures. They often spent as much time working on their file systems as they did developing and maintaining applications software (1:12). The first DBMSs were systems developed by vendors without any real theoretical basis. Of these, only International Business Machines' Information Management System (IMS) remains in successful commercial operation. This can be attributed as much to the large numbers of installations with the software installed as it can to the efficiencies of the system (1:503).

Hierarchical Database Management Systems

Initially, the database systems were hierarchical files systems. In these systems, information is represented in a group of tree-like structures. For each related group of objects, there is one "parent" object at the root of the group (in hierarchical database jargon, the record). All objects (called segments) are connected logically and physically in parent-child relationships. Most importantly, every segment except the root segment, can only be reached after finding all of its parent segments (reference Figure 1). Applications requirements

Structure Diagram



Occurrence Diagram

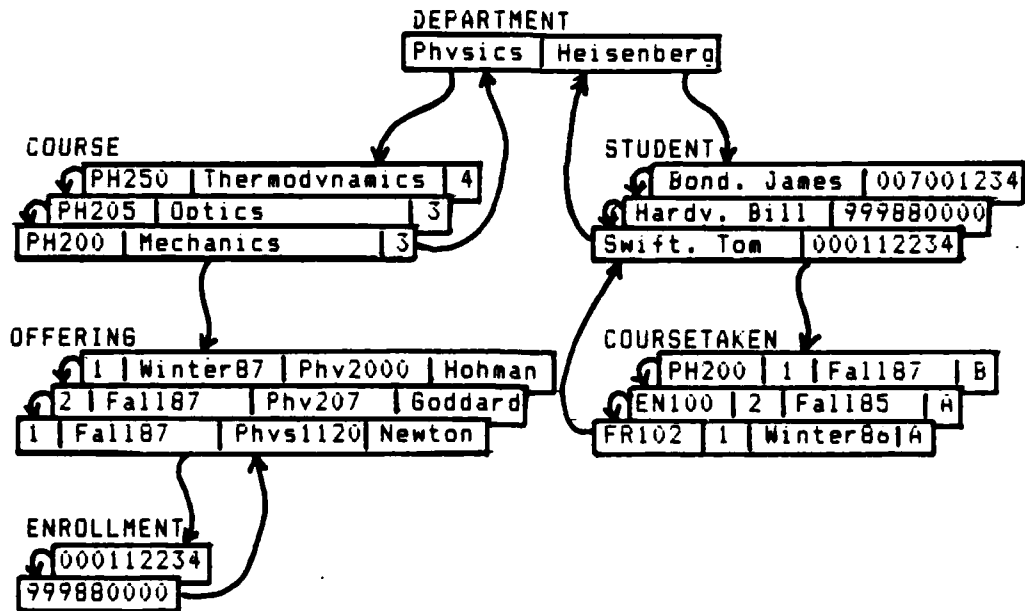


FIGURE 1. HIERARCHICAL DATABASE (1:504)

rapidly outstripped the capability of this "pure" hierarchical data structure, and vendors responded to user requirements with additional data structures to simplify and speed up access to data in hierarchical databases. Today IMS has the ability to partition database records so that applications may reach certain intermediate "roots" directly, and index structures allow preselected types of child segments to be located based on the value of one or more data items within that segment (1:529-533). Other vendors hierarchical systems were similarly modified; Burroughs Corporation's DMS-II uses tree-like and bit-mapped indexing structures, linked-lists, and allowed programmer-defined pointers to link objects in the database in arbitrary patterns (21:4-85, 4-137 - 4-138).

Network Database Management Systems

By the late 1960s, vendors had so modified and diversified their database management systems that the art of software development using database systems was becoming chaotic. The gains in portability and maintainability made with the standardization of most business software around the Common Business-Oriented Language (COBOL) standard was being lost in the growing confusion of differing database management systems. The American National Standards Institute, a joint commercial and governmental standards group had within it the Committee on Data Systems and Languages (CODASYL). CODASYL had been successful in the past in developing standards for the COBOL and other programming languages. In 1969, CODASYL formed the Data Base Task Group (DBTG), which set out to develop a single, agreed-upon industry standard for database management systems. The US government, academia, and major computer vendors were all represented on the group. They developed a specification for a new type of database management system, one which was a significant enhancement of the capabilities of the existing hierarchical systems (13:215-216).

The new standard, the network database, had several new features. While data could still be represented in hierarchies, objects in a database no longer were restricted to a single type of parent object (in DBTG jargon, an owner record). Many different structures could exist in the same database, and access to various types of objects could be improved by specifying the general method used to find and store each different type of record in a database (reference Figure 2). The difference in database structure is significant since several fast access paths could now be built into a database instead of only one. In addition, maintenance of the database was improved by completely separating, for the first time, the logical specification (what record types there were, and how they were connected) from the physical specification (how large disk files would be, how many records could fit on a database, and so on) (1:541-548).

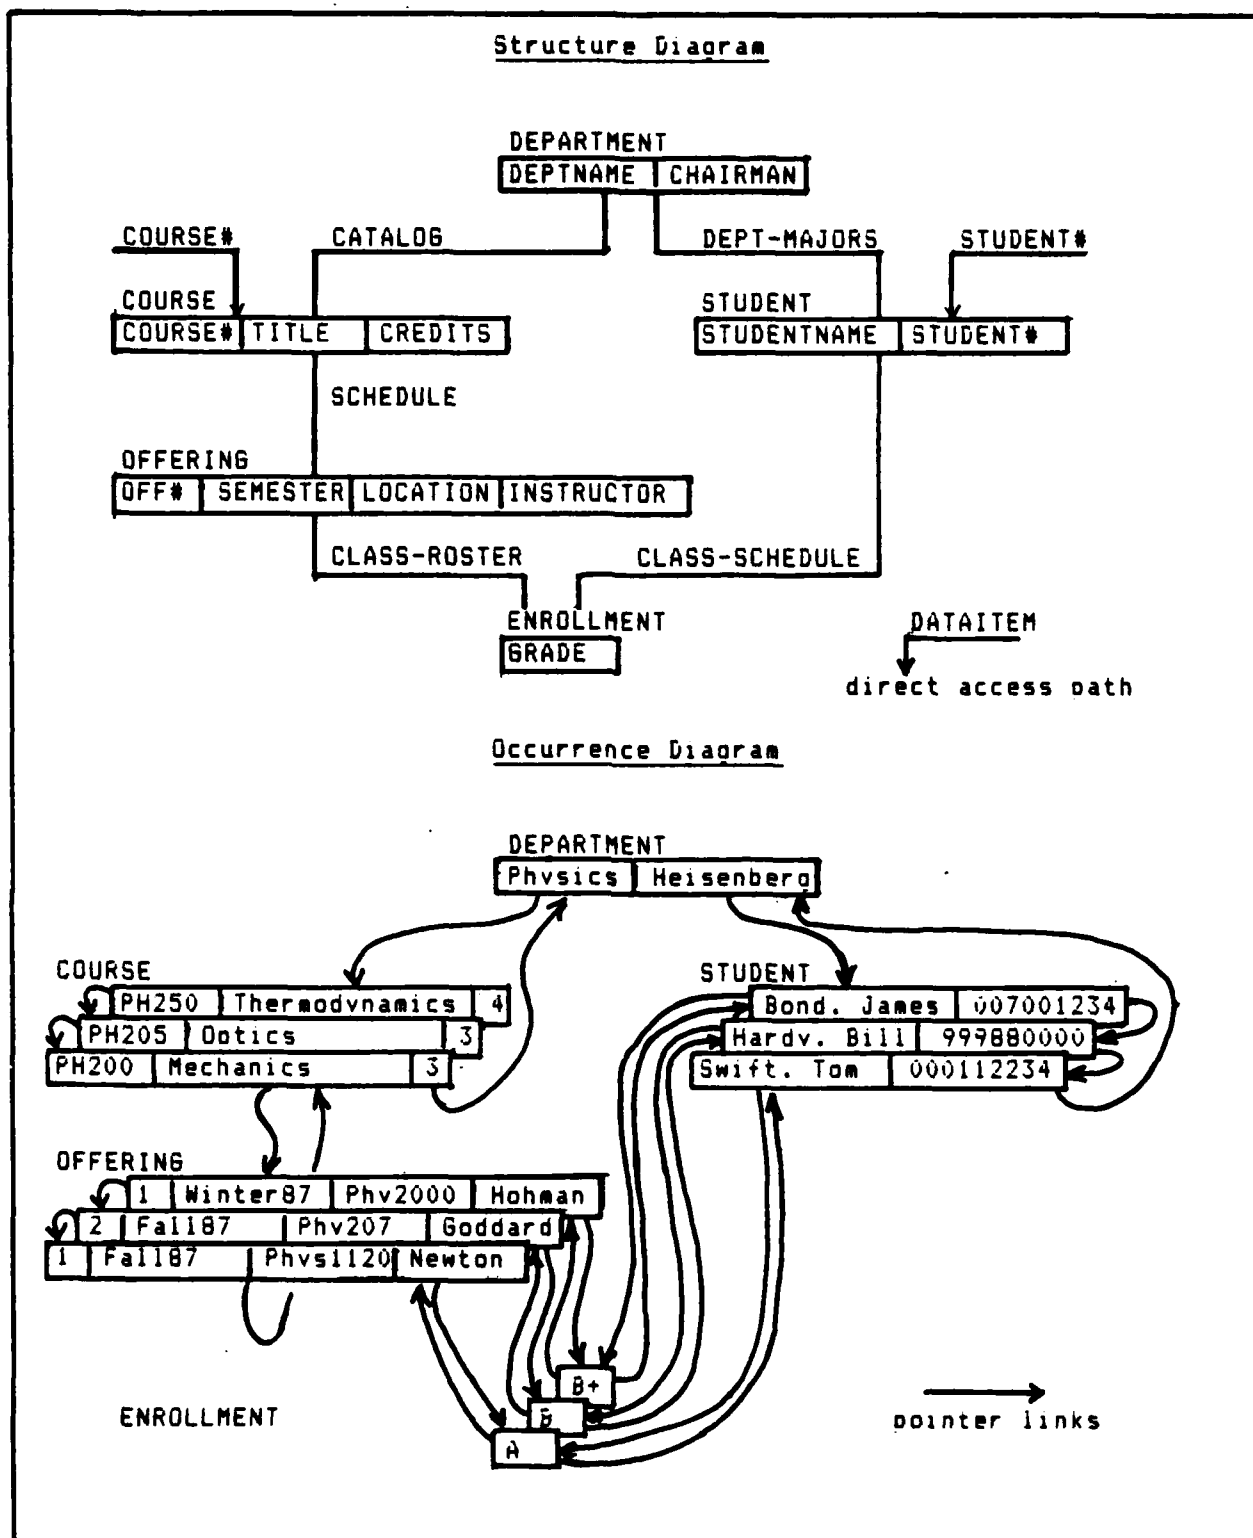


FIGURE 2. NETWORK DATABASE (1:546)

Like the hierarchical systems before it, the "pure" network model has been modified by various vendors to enhance its utility and speed. Today, most network systems (notably Honeywell Corporation's IDS2 and Cullinet's IDMS) offer treelike indexes and pointer arrays to augment data structuring and improve performance (1:561-568).

The Relational Database Model

Even as network database management systems were being developed for commercial use in the early 1970s, their replacement was being created. In 1970, E.F. Codd published his seminal paper on relational database theory (7:--). Codd felt that databases were entirely too ad hoc, and sought to develop a theoretical foundation which would furnish useful tools to database designers, programmers, and users, without being needlessly complicated. Using the concept of the relation, a simple table of data, he demonstrated that this structure could, in theory, satisfy the logical requirements of any database application. Further work by Date at IBM San Jose and his now classic book on database systems (1:--) popularized the concept.

The relational approach abandons the idea of elaborate structures for data in favor of the simple tabular representation. Every "object" in the database is simply a "line entry" or row in one of many tables or relations comprising the database. If the rows in different tables are logically related, they are linked together with matching data values in corresponding fields, called columns (reference Figure 3). The lack of complicated physical data structure is matched with a very simple programming language; this combination makes relational database management systems much easier to learn and use, for both programmers and users (14:217).

The relational database architecture is attractive for several reasons. Simplicity of use and operation is the major reason, but certainly not the only one. The ready availability of third-party applications and productivity tools is another big benefit. (9:144) Very recently, the advent of high-performance relational systems, such as IBM's DB2, has made them attractive to "top end" data processors for whom sheer power was an overwhelming requirement (8:--).

Structure Diagram

DEPARTMENT: (DEPTNAME, CHAIRMAN)

COURSE: (DEPTNAME, COURSE#, TITLE, CREDITS)

OFFERING: (COURSE#, OFF#, SEMESTER, LOCATION, INSTRUCTOR)

STUDENT: (STUDENTNAME, STUDENT#, MAJORDEPT)

ENROLLMENT: (COURSE#, OFF#, SEMESTER, STUDENT#, GRADE)

Occurrence Diagram

DEPARTMENT

DEPTNAME	CHAIRMAN
Physics	Heisenberg
Chemistry	Pauling
English	Strunk

STUDENT

STUDENTNAME	STUDENT#	MAJORDEPT
Swift, Tom	000112234	Physics
Hardy, Bill	999880000	Physics
Bond, James	007001234	Physics
Prince, Charles	777008888	Business

COURSE

DEPTNAME	COURSE#	TITLE	CREDITS
Physics	PH201	Mechanics	3
Physics	PH205	Optics	3
Physics	PH250	Thermodynamics	4
English	EN100	Composition	3

OFFERING

COURSE#	SEMESTER	OFF#	LOCATION	INSTRUCTOR
PH201	Fall87	1	Phvs1120	Newton
PH201	Fall87	2	Chem207	Goddard
PH201	Winter87	1	Phvs2000	Hohman
EN100	Summer86	1	Hum1000	Bacon

ENROLLMENT

COURSE#	SEMESTER	OFF#	STUDENT#	GRADE
PH201	Fall87	1	000112234	A
PH201	Fall87	2	999880000	B
PH201	Winter87	1	007001234	B+
EN100	Summer87	1	777008888	C

FIGURE 3. RELATIONAL DATABASE STRUCTURE (1:124)

DATABASE MACHINES

Development

Paralleling the evolution of database software in the late 1970s, it began to become apparent that traditional computer architectures were ill-suited for database management tasks. The traditional, or "von Neumann," computer architecture is primarily an arithmetic processor, operating on a single "word" of data at a time (usually a number, hence the term arithmetic logic unit, or ALU). Database management on the other hand, often requires that many "words" of data be examined at a time, often in complex ways (3:1-2). For example: How many rated captains with more than six years of commissioned service currently have less than two years on station? What are their dependents' names? At the same time, new concepts in computer hardware were beginning to demonstrate significant performance improvements in arithmetic processing. Pipelined architectures are faster because computer instructions are executed in an "assembly line" fashion; different parts of different instructions are executed simultaneously (4:1145-1148). Parallel architectures are faster because several complete computer instructions are executed simultaneously (4:1100-1104). Researchers began looking at the idea of specially configured or designed computers for database management.

In both approaches, the database services were removed from a centralized computer and placed on a "backend" processor. This backend processor communicates with the mainframe using an agreed-upon set of messages transmitted over a high speed data link. Whenever an applications program needs to use a database, instead of directly invoking the DBMS software, it simply sends a message to the backend processor. The backend, then, receives the message, performs the requested function, and sends an acknowledgement or answer back to the waiting application program on the mainframe (3:13-14).

The Different Approaches

Efforts fell into two main categories, the software and hardware approaches. Software backends are standard, "off-the-shelf" computer systems running specially "tuned" database management and operating system software to deliver high performance. Examples on the commercial market are the Bell Laboratories XDMS system and Britton-Lee's Intelligent Data Management System (IDMS). Other developers felt specialized hardware was necessary to achieve high performance; this approach has led to the hardware backends. These computers typically used either parallel or pipelined architectures to allow the computer to process several records or parts of records at the same time, thereby achieving greatly improved performance through parallelism. Examples on today's commercial market include Intel Corporation's Intelligent Database Processor (IDBP [sic]) and Teradata's Database Computer (DBC) (3:318-319).

Because of the simplicity of the relational data model, most database machines use it as their underlying data model (reference Table 1). However, the hierarchical and network models are also supported by Intel Corporation's iDBP (12:52-53).

Advantages of Database Machines

The database machine approach offers several distinct advantages over conventional general-purpose computers. These advantages fall into three main categories: ease of expansion, improved performance, and low cost.

For some installations, it is much easier to expand the capacity of the computer system by adding a database machine than to replace the entire computer. This expansion can have a dramatic effect, as in the case of one installation using the Teradata DBC/1012, which saw over 99% of its data manipulation workload moved from the mainframe computer to the backend database machine (10:54). Along with this capacity expansion can come a significant performance improvement.

Database machines are specialized processors, providing very fast database operations. Many system developers like the flexibility of the relational database model, but dislike the poor performance of these same systems. By offering a machine designed around the relational data model, vendors have made relational databases efficient to use (9:139).

Perhaps most important, database machines are relatively cheap. The processing power of a database machine is on the order of one-tenth to one-fourth the price of that of a mainframe computer. Teradata Corporation compares one of its larger systems, a 60-processor version of the DBC/1012, priced at \$1.7 million, to the IBM 3084Q, costing \$6.2 million (10:63).

Disadvantages

In spite of all the excitement about database machines, there are some disadvantages to them as well. The major ones are software compatibility, communications overhead, and equipment maintenance.

Software compatibility is the single largest problem in the minds of data processing managers. Older production programs often have been built with file or database structures other than the relational model, and conversion is usually a tedious, expensive, and error-prone process. Some data processing managers simply accept the incompatibility as a necessary evil. They purchase database machines for their new applications in the hope that the older, incompatible software will become obsolete and "whither away" (13:87).

<u>Vendor</u>	<u>Name</u>	<u>data model</u>	<u>capacity</u>	<u>host system</u>
Amperif Corp	RDM 1100	relational	2-16 drives	Sperry 1100, IBM
Britton-Lee, Inc	IDM 500	relational	1-16 drives	IBM
HDR Svstems. Inc	Noah I	relational	1-16 drives	standalone
	Noah II	relational	1-16 drives	VAX 11/750, PC-DOS, Onyx
Intel Corp	iDIS 735	relational	1 40Mb drive	IBM, CDC, Sperry
	iDBP 86/440	relational, network, or hierarchical	1-4 drives	IBM, DEC
Mega/Net Corp	Mega/Net I, II, and III	relational	1-16 drives	Ethernet, LAN, or X.25 networks
Teradata Corp	DBC/1012	relational	2-2000 drives	IBM, Honeywell

TABLE 1. COMMERCIAL DATABASE MACHINES (20:3-6)

Communications overhead has long been a concern of database theoreticians. Date felt that database machines were inherently limited by the communications channels connecting the mainframe to the backend machine (2:348-359). The system which best deals with this concern is the Teradata DBC/1012, which uses a specially designed communications network called a Ynet to eliminate all but the heaviest communications traffic (12:48-50; 17:v. 7-6 - 7-8).

Finally, equipment maintenance can be a problem. Introducing a database machine into a data processing facility adds more points of failure, and on equipment dissimilar to that already installed. Commercial database machine manufacturers are dealing with this problem with a variety of approaches. They use reliable, off-the-shelf components; processors are built to be fault-tolerant; and systems are built in modules which can be repaired without shutting down an entire system (12:53-54, 17:iv).

Chapter 4

COMPARISON OF DATABASE MACHINE AGAINST A CONVENTIONAL DATABASE MANAGEMENT SYSTEM

SELECTING A DATABASE MACHINE FOR THE PDS

The current PDS databases occupy some 84 gigabytes of storage on 80 physical disk drives. The software using this data runs on four Honeywell DPS-8/70 mainframe computers connected to these disk drives (18:Aug 86 - Atch 1). Any database machine useful to the PDS must have the disk capacity to hold the PDS databases and be hardware and software compatible with the existing PDS mainframe computers. As the data in Table 1 indicates, the Teradata DBC/1012 is the only commercially available system which meets these basic criteria.

Relatively few DBC/1012s are installed today: almost all are at IBM installations. However, the Honeywell Corporation is working with Teradata Corporation to develop a connection between the Honeywell GCOS-8 operating system (used by Honeywell mainframes) and the DBC/1012. This capability is expected to be available sometime in 1987 (21:--). Because of the small installed base of DBC/1012s and the absence of commercial Honeywell users, performance data is somewhat limited. This comparison will therefore evaluate the system on qualitative as well as quantitative factors, in an attempt to examine all possible aspects of system performance.

QUALITATIVE FACTORS

Reliability

The DBC/1012 has several features which make it very reliable. Most obvious is the "fallback" capability for "mirroring" valuable data. For those parts of the database where fallback is specified by the database administrator, all the data is duplicated on separate disk storage units (DSUs) within the system. Should a DSU fail, all of its primary data (it, too, contains mirrored data) would be unavailable, but the backup copy of the data would be present, and would immediately be available for use. This allows vital data and applications to operate virtually uninterrupted. The penalty for this level of reliability is, of course, that twice as much disk space is required for protected data.

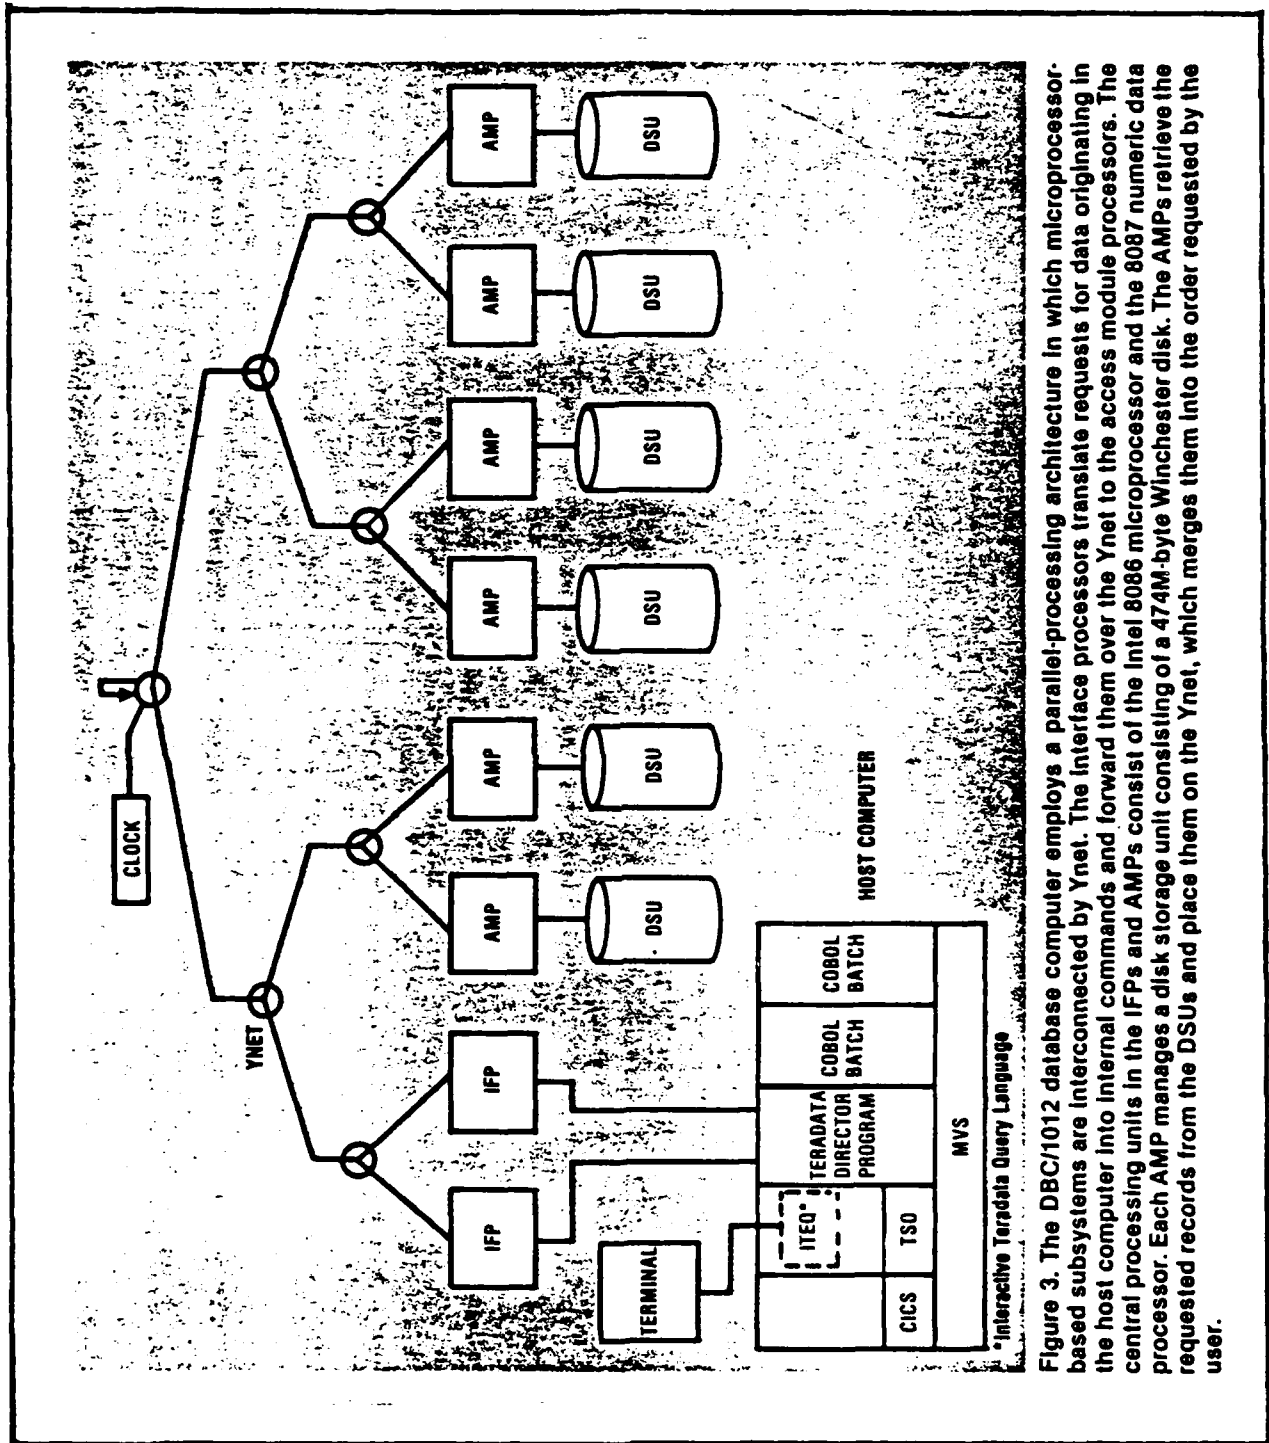


Figure 3. The DBC/1012 database computer employs a parallel-processing architecture in which microprocessor-based subsystems are interconnected by Ynet. The interface processors translate requests for data originating in the host computer into internal commands and forward them over the Ynet to the access module processors. The central processing units in the IFPs and AMPs consist of the Intel 8086 microprocessor and the 8087 numeric data processor. Each AMP manages a disk storage unit consisting of a 474M-byte Winchester disk. The AMPs retrieve the requested records from the DSUs and place them on the Ynet, which merges them into the order requested by the user.

FIGURE 4. TERADATA DBC/1012 CONFIGURATION (12:49)

In addition to the fallback protection, there are several other features which add to the system's reliability. The DBC/1012 is built almost entirely with off-the-shelf components, such as the Intel 8086 processor and the Winchester disk drive (12:53-54). All these parts are high-performance, very reliable, and relatively inexpensive. The system is built in modules, with the Vnet communications link serving as the only connection between the various processor units (reference Figure 4). This "loose connection" allows an ailing processor to be disconnected from the system and repaired or replaced while the rest of the computer continues to operate (12:54). The vnet itself is a dual-channel communications link, which can continue operating at a reduced rate even if one channel should fail (17:7-8). In addition, the Vnet can furnish data to the user in sorted order, independent of the ordering of the data on the disk storage units (12:47-49).

Capacity

Besides reliability, the DBC/1012 provides excellent capacity, both for the workload it is originally acquired for, and for future expansion. The system can be configured to store up to 2.1 terabytes or 2150 gigabytes, the equivalent of 8600 bytes of data for every person in the United States or 430 bytes of data for every person on Earth! Because the disk space can be added in increments as small as 474 megabytes, it is relatively easy to build a system which is exactly the right size for a given application. As processing or storage requirements increase, processing units (IFPs and AMPs, reference Figure 4) can be added to maintain or improve performance, and the system will automatically reorganize its databases for peak performance, without complicated human intervention. Furthermore, IFPs can be added to the system to connect it to multiple mainframe computers, allowing more "front end" power to be added, more terminal users supported, and different types of computers to share the same databases. (19:--).

QUANTITATIVE FACTORS

Performance

Besides the qualitative, "nice to have" features, the DBC/1012 offers a distinct performance advantage for large database applications. The performance gains for the DBC/1012 vary according to the type of application, but overall are very impressive. Batch update and retrieval programs especially show significant gains. To estimate these gains, the performance of the current Active Airman Master Personnel File system will be compared with the projected performance using the DBC/1012.

The Active Airman (AA) system uses a set of software called the Generalized Update System (GUS). GUS is a database and software system

used for maintaining the master personnel files and other sets of information with similar structure. It is designed to efficiently maintain the data on a large number of similar individual records, for example all Air Force officers, all Air National Guard airmen, or all Air Force suggestions. At its heart is a batch program which processes both update and retrieval transactions against a specific master personnel file. It consists of a basic "host" program which reads and writes the database. The host is linked to many transaction "modules," smaller subprograms which contain the logic for executing each separate type of transaction. Several reports are generated from the outputs of the batch update program, based on output transactions which are written to files as part of the batch update process.

Once updated, the GUS databases are used for a wide variety of data retrievals. Standardized, periodic report programs are run for many users who have recurring needs for information. One-time reports can be generated using the Air Force-developed ATLAS retrieval language, which allows a personnel specialist to specify his information requirement and report format, and run it as scheduled production program. The ATLAS retrieval capability accounts for some 53% of all computer use in the headquarters-level PDS (18:--).

Typical performance times for GUS processes supporting the AA master file are given in Table 2, along with estimated times for performing similar functions using the DBC/1012. The DBC/1012 times were derived from the vendor-supplied performance curves for transaction rates based on the number of processors in the system, and assuming that the system was configured with 6, 10, or 20 AMPs (reference Figure 5). As can be seen, the DBC/1012, even in a minimal configuration, can readily support the functions of the PDS.

Price

Lightning performance is always desirable, but it does no good if it is prohibitively expensive. Surprisingly, the DBC/1012 is a relatively inexpensive system. A study by the Rome Air Development Center showed that even small DBC/1012 systems compare favorably with conventional general-purpose systems with similar capacities (17:7-14 - 7-15). For a system large enough to support the PDS (20 gigabytes or larger), the cost of the system is within an order of magnitude of the cost of disk storage alone. The current storage required for the database portions of the PDS, some 84 gigabytes, requires 80 MSU501 disk storage units, costing approximately \$50,000 each, or \$47,620 per gigabyte of storage (23:--). While the price of a Teradata system with this capacity is not available, the largest system with a published price (4 IFPs, 20 AMPs, and 40 DSUs) has a total system cost of \$1,475,000, or about \$71,600 per gigabyte of storage (reference Table 4). While this cost is 50% greater than the cost of conventional storage, it ignores the 16 million instructions per second (MIPS) of processing power added to the system, and the workload relieved from the existing mainframe.





FIGURE 5. TERADATA DBC/1012 PERFORMANCE CURVES (19:27)

process	number of rows	run times (processor/elapsed) in hours			
		DBC/1012 estimates			
		current	6 AMP	10 AMP	20 AMP
AA update	100.000 - 150.000	5.8/20.0	2.8/9.6	1.7/5.9	0.9/3.0
ATLAS retrieval	540.000	2.4/2.6	0.13/0.17	0.07/0.10	0.04/0.05
monthly extracts	12.000.000	8.7/30.2	2.8/9.5 *	1.6/5.6 *	0.8/2.8 *

* assumes sorting by social security number is performed
by the DBC/1012 in parallel with other processing

TABLE 2. AA PERFORMANCE ESTIMATES (24:--)

IFPs	AMPs	DSUs	disk space	processing power (MIPS)	price	price/Gb
2	2	4	2.1 Gb	4.0	\$ 335.000	\$159.500
4	4	8	4.1	8.0	525.000	128.000
4	8	16	8.2	12.0	770.000	93.900
4	12	24	12.4	16.0	950.000	76.600
8	20	40	20.6	28.0	1.475.000	71.600

TABLE 3. PRICES OF TERADATA DBC/1012 SYSTEMS (19:25)

Chapter 5

CONCLUSIONS, FINDINGS AND RECOMMENDATIONS

INTRODUCTION

Database machines, and the Teradata DBC/1012 in particular, have several distinct advantages over conventional general-purpose computers which recommend them. The most significant are their cost, capacity, performance, and reliability.

CONCLUSIONS

Cost

Considering its capabilities, the DBC/1012 is a bargain. Current disk storage, using Honeywell MSU501 disk drives costs approximately \$50,000 per gigabyte. Disk storage using the DBC/1012 costs \$71,000 per gigabyte in the size range needed for the PDS. While this is half again as much as the cost of MSU501 storage, it includes the inherent processing power of the DBC/1012 (16 - 28 MIPS), and the Teradata DBMS.

Capacity

A relatively modest configuration of the DBC/1012 would provide sufficient storage capacity for the entire PDS database, even with complete data redundancy using the fallback method of duplicating data. Both performance and storage capacity can be easily increased by adding hardware modules (IFPs and AMPs), rather than making large-scale modifications or replacements of mainframe computers. The added processing power of the DBC relieves the mainframes of a considerable workload, thereby extending its useful service life. Additionally, performance of the DBC can be improved through hardware upgrades to the DBC modules themselves, by upgrading the processors, adding memory to the hardware modules, or by increasing the capacity of the disk units. Finally, the DBC can be connected to several "front-end" mainframe computers, making data sharing among various computers very easy.

FINDINGS

Performance

The DBC/1012 is one of the few systems which does not rapidly fall victim to the law of "diminishing returns." In the range of systems large enough to operate the PDS, the performance of the DBC increases linearly as the system is expanded. This means that most foreseeable performance improvements can be made simply by adding hardware.

Batch update performance is very sensitive to the number of AMPs available in the system, and is severely constrained in small systems. This is not surprising, considering that each transaction deals with a single record, and once every AMP is working on a transaction, any subsequent transaction will have to wait for the AMP processing its "target" record to complete its transaction before proceeding. As long as batch updates are performed a "record-at-a-time," this will be a major bottleneck for batch updates. Transactions which could be "broadcast" to the entire database to update all applicable records would greatly increase the performance of batch update programs. While this is certainly a desirable situation, the real world may not be as cooperative in making these "mass update" transactions the predominant way of doing business.

Batch inquiry performance is the area where impressive gains can readily be realized using the database machine. Row retrievals are very fast, and the Ynet's sort capability is a processing bonus which is conceivably as important as the database management facility. Because of the great performance gains which can be realized from being presented with sorted data by the database machine, any facility using the DBC/1012 for large databases should procure the AMPs with the maximum amount of sorting capacity available.

Reliability

The DBC/1012 is highly reliable for three basic reasons: it allows complete redundancy in data storage, it is constructed of proven, off-the-shelf components, and its configuration is highly modular. The data redundancy and modularity are especially significant, since they allow components of the system to be repaired while the rest of the system continues operating. This combination prevents all but the most catastrophic of failures (such as a total power outage) from putting the system out of service.

Conversion

Converting software to run on the DBC/1012 requires that the software use relational databases. This is easiest in those systems

with simple data structures and small programs which are not tightly bound to the structure of the data. It is most difficult where the data is highly structured, the programs are large, and program logic is tightly bound to the data structure. Unfortunately, the larger GUS systems of the PDS fall into the latter category. To some extent, this contributed to the difficulty experienced in converting these systems during the REACQ project. This will be a problem in any future conversion involving a departure from the current batch, record-at-time processing concept.

RECOMMENDATIONS

The database machine offers capabilities that the PDS cannot afford to do without. Operating on a database machine, the system will enjoy high performance, easy capacity management and improvement, and high reliability, all at a very reasonable cost. The problem is getting from here to there: determining the actual system size and converting the applications software are two major hurdles to cross.

System Sizing

Sizing a system is important for a simple reason: if the hardware does not have sufficient storage or processing capacity for a set of applications, no amount of software wizardry will make the system perform. Fortunately, sizing the DBC/1012 system is relatively straightforward. An initial estimate of system size can be made based on the size of the database to be supported and the number of mainframe computers to be connected to the database machine. These estimates can then be used as "first guesses" that can be further refined using the Rome Air Development Center model for performance estimation. (17:vi) Using this model should give a very close estimate of the necessary system size.

Software Conversion

Software conversion was the major difficulty of the REACQ project, and it could conceivably make installing a database machine impractical. Because benchmarking is an integral part of selecting and acquiring a hardware system, the Directorate of Personnel Data Systems needs to begin prototype development of a relational database implementation as soon as possible. The prototype effort should have three primary goals in order to be useful: concept development, software methodology, and benchmarking.

Concept development is simply deciding on basic overall processing strategies, features to take advantage of, and pitfalls to avoid. It should provide the basic "road map" for the prototype developers to keep their efforts consistent with the mission and functions of the PDS and

to avoid duplication or false starts. With a well-thought out concept, the software for the prototype can be developed.

The software methodology is important because it will determine how production software will be developed in an actual database machine environment. This can be begun on the existing Honeywell DPS-8/70 system using the installed Personal Data Query (PDQ) system. This part of the prototype effort will validate the prototype concept and provide necessary "lessons learned" well before the actual database machine is ever used.

The benchmark is the "moment of truth" for selection of the database machine. With a prototype developed using PDQ, conversion of the software to run on the DBC/1012 should be minimal. The benchmark can be compared to the actual production software currently in use, and to the prototype software as it was run on the mainframe system. Given a significant performance improvement using the benchmark trial, the DBC/1012 should be purchased for installation as an integral part of the PDS.

BIBLIOGRAPHY

H. REFERENCES CITED

Books

1. Date, C. J. An Introduction to Database Systems, Vol 1. Reading, Massachusetts: Addison-Wesley Publishing Co., 1985.
2. -----, An Introduction to Database Systems, Vol 2. Reading, Massachusetts: Addison-Wesley Publishing Co., 1985.
3. Ozkarahan, Esen. Database Machines and Database Management. Englewood Cliffs, New Jersey: Prentice-Hall, Inc., 1986.
4. Ralston, Anthony and Reilly, Edwin D., Jr. (editors). Encyclopedia of Computer Science and Engineering. New York: Van Nostrand Reinhold Co., 1983.

Articles and Periodicals

5. "Teradata Speeds Up DB Processing." MIS Week, 13 January 1986, p. 18.
6. "The United States Air Force in Facts and Figures." Air Force Magazine, Vol. 69, No. 5 (May 1986), pp. 181-197.
7. Codd, E.F. "A Relational Model of Data for Large Shared Data Banks." Communications of the ACM, Vol. 13 No. 6 (June 1970), pp. 377-387.
8. Date, C.J. "How Database Systems Perform." Computerworld, (13 February 1984), pp. 10/15-20.
9. Epstein, Robert. "Why Database Machines?" Datamation, Vol. 29 (July 1983), pp. 139-144.
10. Myers, Edith. "Database Machines Take Off." Datamation, vol. 31 (14 May 1985), pp. 52-63.
11. Neches, Philip M. "Hardware Support for Advanced Data Management Systems." Computer, vol. 17 No. 11 (November 1984), pp 29-41.

CONTINUED

12. Shemer, Jack and Neches, Philip. "The Genesis of a Database Computer." Computer. Vol. 17. No. 11 (November 1984). pp. 42-56.
13. Sigler, Lee. "Relating to Database Machines." Datamation. Vol. 32 (1 April 1986). pp. 83-90.
14. Wiederhold, Gjo. "Databases." Computer. Vol. 17 No. 10 (October 1984). pp 211-223.

Official Documents

15. Air Force Manpower and Personnel Center. United States Air Force Personnel Management Information Systems Plan. Randolph Air Force Base, Texas: AFMPC/DPMD. 31 December 1985.
16. Headquarters, US Air Force. AFR 23-33: Organization and Mission - Field. Air Force Manpower and Personnel Center. Washington, DC: HQ USAF/MPCP. July 1979.
17. Rome Air Development Center. RADC-TR-85-273. CSI Teradata Study. Griffiss Air Force Base, New York: RADC/Measurement Concepts Corporation. March 1986.

Unpublished Materials

18. Directorate of Personnel Data Systems. Air Force Military Personnel Center. Randolph Air Force Base, Texas. Monthly computer system performance briefings. October 1985 to October 1986.
19. Teradata Corporation. Los Angeles. "The Database Computer (DBC/1012)." briefing. undated.

Other Sources

20. "All About Database Machines." Datapro Report 70Db-001LB-105. Delran, New Jersey: Datapro Research Corporation. June 1985.

CONTINUED

21. Burroughs Corporation. B6000/B7000 Series DMSII DASDL Reference Manual. Detroit: Burroughs Corporation. March 1978.
22. Engstrom, Walter. Honeywell Large Computer Products Division. Phoenix, Arizona. Letter. 14 August 1986.
23. Essendreis, David. Technical Advisor to the Director of Personnel Data Systems. Air Force Military Personnel Center. Randolph Air Force Base, Texas. Telecon. 5 February 1987.
24. Robidoux, R.A., Capt. USAF. Database Analyst. Database Branch. Air Force Military Personnel Center. Randolph Air Force Base, Texas. Telecon. 3 February 1987.

B. RELATED SOURCES

Books

- Vasta, Joseph A. Understanding Database Management Systems. Belmont, California: Wadsworth Publishing Co., 1985.
- Wiederhold, Gjo. Database Design. New York: McGraw-Hill Book Company, 1977.

Articles and Periodicals

- Codd, E.F. "Relational Database: A Practical Foundation for Productivity." Communications of the ACM. Vol. 25 No. 2 (February 1982), pp. 109-117.
- Stonebraker, Michael. "Operating System Support for Database Management." Communications of the ACM. Vol. 24 No. 7 (July 1981), pp. 412-418.

END

7-87

DTIC